ANTLER

ANTLER INTELLIGENCE | JUNE 2026

What Your Portfolio Thinks About Disruption

A framework for reading innovation exposure and disruption risk across institutional portfolios

AUTHOR: ROSALIND BAZANY, PARTNER

I have spent time on both sides of how markets interpret technological change. Earlier in my career, working across public markets and multiple asset classes, the lens was largely one of risk management, allocation frameworks, and benchmark construction. Early-stage venture capital offers a different vantage point. You sit closer to where companies are being built before they are priced, before they are widely known, and often before the sectors they may eventually reshape have fully recognised the pressure coming their way.

From this side of the market, the more important issue increasingly looks informational rather than return driven.

The market now appears to operate across two parallel layers. One is the disruption layer: the systemic forces actively revising the assumptions on which a significant portion of institutional capital is priced. The other is the innovation layer: the companies being built on top of those forces, mostly private, mostly early-stage, and often outside the institutional capital stack entirely.

Most large portfolios already have meaningful exposure to the first layer and limited exposure to the second. The reasons are understandable: mandate restrictions, governance frameworks designed for later-stage evaluation, and fee structures difficult to justify against public market benchmarks. Investors best positioned to absorb long-duration illiquidity are often the furthest from the layer where the earliest forward signals emerge.

To test whether the observation held up beyond intuition, I built a framework and applied it to publicly disclosed holdings across some of the largest pension funds in developed markets using December 2025 positions.

I am working from public disclosure and not claiming to have a full picture of any fund's actual positioning. PwC's 2025 Global Investor Survey found that 53% of investment professionals describe the companies they invest in as having high or extreme exposure to technological disruption.

The framework

The method maps every material position against two axes.

Disruption Risk measures the degree to which a position depends on business model premises currently being revised by systemic forces. Systemic and cyclical pressures behave differently. Cyclical pressure recovers when conditions normalise; systemic disruption involves a permanent revision to the underlying logic that supported a valuation: the demand assumptions, the competitive barriers, the regulatory framing. Positions in the second category require an explicit view on the pace of that revision, not just its direction.

Innovation Exposure measures the degree to which the portfolio participates in the companies and capital generating the disruption, rather than merely receiving it. This axis has three layers and conflating them is the most common source of false confidence in how institutional portfolios describe their innovation positioning.

- The infrastructure layer covers companies building the computational, physical, and financial substrate on which the next economy operates: hyperscalers, chip manufacturers, semiconductor tooling, and energy infrastructure supplying the capacity for AI workloads. Owning these is a bet on the buildout of a platform, not on what gets built on top of it.
- The application layer covers companies actively deploying that infrastructure to displace incumbent business models, among them AI-native financial services providers compressing traditional banking margins, climate technology companies repricing the energy sector's asset base, and health technology businesses restructuring diagnostics and distribution in ways that revise chronic disease device economics. Most significant players in this layer are private at the stage where the displacement economics are most asymmetric.
- The formation stage covers pre-revenue or early-revenue companies at the point where their eventual displacement of incumbents is speculative, but where the return profile is most asymmetric, and where the earliest signals about which categories are forming are clearest. Most institutional portfolios have very limited exposure to this layer, for reasons I think are understandable: mandate restrictions that explicitly prohibit pre-revenue investing, fee structures that may be difficult to justify against public market benchmarks, governance frameworks not designed to evaluate companies without revenues, and the simple fact that top-quartile managers at this stage are capacity-constrained.

Within the Disruption Risk axis, positions are classified against five vectors. Each represents a distinct mechanism by which business model premises are being revised.

- *AI Displacement:* business models dependent on labour-intensive cognitive functions that AI agents, automation, and decision engines are demonstrably replacing. Retail banking operations, insurance processing, legal and professional services, employment intermediaries. Positions here require a view on margin trajectory over a three-to-seven-year horizon that standard sector analysis is not calibrated to provide.
- *Climate Transition:* operating economics priced on carbon-intensive premises that policy, technology, and shifting capital costs are actively revising. Fossil fuel extraction, coal-dependent infrastructure, conventional automotive, carbon-intensive materials. The vector is asymmetric: risk for those holding premises being revised, and significant opportunity for those positioned on the other side.
- *Biotech and Health Disruption:* disease burden, utilisation, and drug development premises being revised faster than most models anticipated. AI-accelerated drug discovery is compressing the timeline from target identification to clinical candidate. AI diagnostic tools are restructuring clinical workflows and the procedural volumes that underpin medical device and hospital economics. Both are moving faster than current equity valuations in the sector tend to reflect.
- *Regulatory Fragmentation:* global business models built on cross-border data flows, payment rails, and platform economics that assumed regulatory convergence, now facing divergence through data localisation requirements, AI Act compliance variation, sanctions regimes, and trade policy bifurcation.
- *Geopolitical Supply Chain:* production and distribution architectures exposed to fracture lines in semiconductor trade restrictions, critical mineral dependencies, energy supply concentration, tariff regimes, and conflict proximity. This vector cuts across sectors in ways that geography-based analysis tends to miss.

The five vectors do not operate on the same timeline. Geopolitical Supply Chain risk is active in certain positions today. AI Displacement is revising margins across most affected sub-sectors on a medium-term horizon. Climate Transition plays out across decades but carries non-linear policy shock risk that can accelerate repricing abruptly. Understanding which vector is relevant to a given position, and at what velocity, matters as much as identifying it.

Innovation exposure also has a geography. The disruption vectors hitting most large institutional portfolios are global in origin: AI compression of banking margins, energy transition pressure on resource companies, pharmaceutical innovation revising device market economics. The companies being built on the other side of those trades are distributed across San Francisco, Stockholm, London, Singapore, and increasingly across markets in APAC, MENAP, and Continental Europe. Formation-stage capital that is geographically domestic to a given fund's home market may be excellent capital, but it is not necessarily positioned on the same displacement vectors the fund is most concentrated in.

What I found

Applied to publicly disclosed holdings across the cohort, the framework surfaces four observations consistent enough across funds, jurisdictions, and portfolio structures to warrant examination. All are drawn from the visible portion of each portfolio only.

AI infrastructure concentration is coordinated, not diversified

Across the funds looked at, AI infrastructure concentration tends to cluster between 15 and 25% of the total equity book, concentrated in the same eight to ten names: the dominant hyperscalers, the leading semiconductor manufacturer, the principal chip designer driving the AI compute cycle, the key tooling companies in the semiconductor supply chain. When these names appear at consistent weights across multiple institutional equity books, correlation analysis still classifies them as diversified across sectors. They are not. They represent coordinated exposure to one resolution of the AI transition: that value accrues primarily to infrastructure incumbents, that the application layer does not materially compress their margins, and that current capital expenditure cycles are well-calibrated to actual demand.

The leading chip designer appeared as one of the largest single equity positions, at weights of between 3 and 6% of the total equity book. The position is real conviction, and the returns have validated it. The question the framework is designed to ask is a different one: whether owning the infrastructure layer, at valuations that already price in significant AI optionality, constitutes genuine innovation exposure or consensus market exposure to the current cycle's winners. These are concentrated bets on one resolution of the AI transition.

A further observation on the semiconductor positions specifically: the largest chip manufacturer in most portfolios spans two disruption vectors simultaneously, functioning as a near-pure infrastructure-layer bet and one of the highest concentrations of geopolitical supply chain risk in the same book. Most advanced AI chips driving the current cycle are manufactured in a single jurisdiction. Standard sector analysis does not surface both dimensions together, and most risk reporting does not either.

Banking sector exposure is roughly double what equity reporting shows

The banking sector exposure visible in the equity book consistently understated the actual position once fixed income holdings in those institutions were aggregated. The largest domestic and global banks appear as meaningful equity positions; they then reappear across the fixed income book in bonds at comparable scale. One name, the same AI Displacement vector, in a separate part of the portfolio report. Asset-class-level reporting is designed to show diversification across buckets, not concentration within a single disruption vector across them.

The practical implication: an investment committee reviewing banking sector disruption risk from equity-only analysis is looking at roughly half the actual exposure. The fixed income

component compounds the equity concentration rather than offsetting it and adds credit and counterparty dimensions that sector analysis is not built to capture. For funds where domestic banks are simultaneously the largest equity positions, significant fixed income issuers, and primary cash counterparties, the cross-asset aggregation produces a number that is materially larger than any single-asset-class view would suggest.

Drug innovation and device exposure are reported as separate positions in one trade

Across the cohort, funds hold positions in the leading GLP-1 drug manufacturers alongside meaningful positions in the global leaders in chronic disease device categories whose volumes are being revised by those same drugs. The two positions sit in different sector classifications, in some cases in different asset classes, and are reported in ways that make them appear as unrelated exposures. A disruption lens reads them as a single trade: a long position in a pharmaceutical innovation cycle and a long position in the device market assumptions that cycle is actively revising.

The GLP-1 dynamic is the most visible current instance of this pattern, but the underlying mechanism is not specific to it. AI-accelerated drug discovery compressing the timelines to clinical candidates across multiple therapeutic categories, diagnostic tools restructuring the procedural volumes that underpin medical device economics, pharmaceutical innovation revising the chronic disease burden assumptions on which insurance and device businesses are priced: the same cross-sector revision dynamic recurs across the health book in ways that sector-level reporting does not aggregate. The pattern is not a risk management failure; standard classification frameworks were not built to show it.

Formation-stage exposure is limited or absent across most portfolios

Across the funds I examined, the pattern holds regardless of fund size, total alternatives allocation, or the sophistication of the private markets programme: exposure to early-stage venture capital is either absent from the disclosure or limited in scope and usefulness. The private and alternative books are professionally managed and, in several cases, distinguished. The gap is in access to the innovation signal, not in the quality of what surrounds it.

The geographic pattern compounds this. Where early-stage venture exposure exists, it tends to be domestic, concentrated in a small number of manager relationships, and calibrated more to benchmarking logic than to the geography of disruption. The displacement vectors hitting the largest holdings in most portfolios are not playing out domestically first. AI compression of banking margins, energy transition pressure on resources companies, software-led transformation of enterprise services: these are being built in San Francisco, Stockholm, London, and Singapore. Domestic early-stage capital hedges against domestic displacement risk. For portfolios whose largest disruption concentrations are global, the geographic mismatch is material.

The categories forming at pre-seed and seed are the most direct leading indicator of where competitive pressure on existing holdings is building. Public market pricing captures that signal after it has moved. The signal appears upstream first.

Three questions sit underneath the framework

These are observations from public disclosure, not a full audit of any fund's positioning. Three questions keep coming back to me when I look at portfolios through this lens.

1. The calibration question. Where in the portfolio is there a long position in business model premises currently being revised by systemic forces, and is that a conscious, priced-in view or an inherited position that has not been reassessed against the velocity of the revision it faces? Not sectors in broad terms, but specific premises: that branch-based banking retains pricing power through an AI-enabled challenger cycle; that chronic disease device volumes are durable against pharmaceutical innovation revising their end markets; that semiconductor infrastructure at current valuations represents innovation access rather than consensus market exposure; that global platform economics persist as regulatory convergence fragments. Investment committees tend to have views on which sectors face disruption. They less often have documented positions on the speed.

2. The position question. Across the total portfolio including private markets, how much capital sits at the earliest stage of the displacement vectors the portfolio is most concentrated in, and is it geographically matched to where those displacements are being built? For most portfolios the honest answer reveals a gap that qualitative investment discussions tend not to expose. The early-stage exposure that exists is typically small, predominantly domestic, and calibrated more to benchmarking logic than to displacement geography.

3. The horizon question. For each disruption vector that is material in the portfolio, what is the fund's view on the pace of repricing, and does the current position reflect that view or is it inherited from a prior allocation decision that assumed a different velocity? A position defensible under a seven-year transition horizon carries different risk if that transition is compressing to three. The pace question is the one most consistently absent from investment committee discussions about disruption.

The price discovery gap

The most valuable thing early-stage investing offers large institutional portfolios looks informational rather than return-driven. Across full-year 2025, 136,000 founders applied to build with Antler globally. We backed 394 of them, roughly one in 346, approximately twelve times more selective than Harvard. What those founders choose to work on is often a leading indicator of where competitive pressure is building long before it becomes consensus.

Consumer AI dominates the headlines. Yet consumer technology's share of Antler's annual investments fell from 27% in 2022 to 15% in 2025. The venture layer signal has been moving elsewhere: toward enterprise infrastructure, applied AI at the sector level, physical AI, and deep technology categories that remain largely absent from later-stage deal flow. The number of physical AI companies in Antler's annual cohorts was zero in 2022, six in 2023, fifteen in 2024, and twenty-five in 2025. That is a category forming in real time, twelve to eighteen months before it begins appearing meaningfully in growth-stage valuations. Portfolios with no visibility into this layer will likely encounter it only once it starts repricing the industrial and hardware positions they already hold.

Quantum computing, physical AI, applied science at seed: these are not abstract future themes. They are where capable founders are choosing to build across twenty-six markets, filtered through one of the most selective founder funnels in early-stage investing.

The deeper irony is the mismatch for long-duration investors. Pension funds, sovereign wealth funds, and endowments are structurally among the best positioned to absorb illiquidity, hold through J-curves, and participate in long-horizon technological shifts. Yet many portfolios remain concentrated in assets where information about disruption arrives only once repricing has already begun.

The pattern recurred across the portfolios analysed, driven less by idiosyncratic investment decisions than by institutional incentive structures that make tracking-error minimisation easier to maintain than disruption positioning. The gap is visible from public disclosure alone, which ultimately prompted the development of a simplified self-assessment version of the framework. The more important question it raises is not whether every institutional portfolio should allocate to early-stage venture capital. It is whether portfolios exposed to systemic technological change can afford to remain structurally distant from the layer where the earliest signals about that change first emerge.

ABOUT ANTLER INTELLIGENCE

Antler is one of the most active early-stage venture firms globally, with US$1.3B+ AUM, operations across 26 markets, and a portfolio of 1,900+ companies spanning every major technology market. We invest from pre-seed and seed, where much of our portfolio sits, through to Series C via regional early-stage funds and global follow-on vehicles. That structure gives us something few investors have: a longitudinal, cross-geography dataset of founder behaviour, category formation, and technology adoption that begins before most institutional investors are in the room.

Antler Intelligence is our research and market insight function designed for an LP audience. The views expressed here are based on publicly available disclosure documents and Antler's proprietary analytical model. They represent a research perspective on portfolio construction dynamics, not investment recommendations or the position of any Antler fund. We publish this analysis because the early-stage vantage point is underrepresented in institutional market discourse, and because our LPs and ecosystem partners benefit from understanding what we are seeing before the market prices it in.

DATA AND DISCLOSURE

This article may contain forward-looking statements based on current assumptions, which are subject to change. No obligation is undertaken to update this material. It is not intended for distribution in any jurisdiction where such distribution would be contrary to applicable law.

Holdings data referenced in this analysis is sourced from publicly available position-level portfolio disclosures published by institutional funds in accordance with their regulatory requirements. Data reflects December 2025 portfolio positions. Funds were selected based on disclosure quality and scale; no fund is named in this analysis. All observations are stated as cohort-level patterns, not attributions to specific institutions. Externally managed mandates hold underlying securities not visible in direct disclosure; actual exposure to individual sectors, issuers, and asset classes will differ from disclosed positions. This analysis should not be taken as a complete characterisation of any fund's investment approach, mandate, or performance.

The disruption vector model applied here is a qualitative analytical lens developed by Antler, not a quantitative risk model. Vector classifications reflect Antler's assessment of the systemic forces affecting relevant business models at time of writing; reasonable analysts may reach different classifications.

Antler operates in the formation-stage market that this analysis identifies as underrepresented in most institutional portfolios. Readers should weigh the analysis with that commercial interest in mind. This document does not constitute investment advice, a personal recommendation, or an offer to invest in any Antler fund or vehicle.

FOR INVESTMENT TEAMS INTERESTED IN APPLYING THIS APPROACH

A self-assessment diagnostic that produces an indicative read of innovation access and disruption exposure at portfolio level is available for illustrative purposes. Twenty questions, approximately ten minutes. For a conversation, contact intelligence@antler.co.